



Optoelectronic Intelligent RAM Multiprocessor: *Opto-IRAM*

Objectives:


- Scalable, high-performance, parallel processing with general-purpose utility
- Low processor–memory latency
- Low power-delay-volume product
- Exploit VLSI Photonics core technologies

Approach:

- IRAM microprocessor architecture for low local memory latency
- Simple cache-coherent protocol and wide optical transport for low inter-chip latency
- High-performance scalar processor plus massively parallel Computational RAM (C-RAM) processing element array
- Design automation to optimize electrical, optical interconnect resource utilization
- Advanced free-space OE module concept

Accomplishments:

- OCRAM architecture
- OE module physical design for virtual prototyping
- Placement and routing tool for mixed-mode interconnect optimization



Optoelectronic Intelligent RAM

DARPA VLSI Photonics Review, July 30, 2001





Program Team

Georgia Tech ECE/MiRC:

- **T. J. Drabik (PI)**
- **A. Chatterjee (PI)**
- **J. D. Meindl (PI)**
- **S. Yalamanchili (PI)**
- **R. C. Copeland (GRA)**
- **C.-S. Seo (GRA)**

Sun Microsystems:

- **H. Davidson**
- **R. Lytel**
- **A. Saulsbury**

Optoelectronic Intelligent RAM

DARPA VLSI Photonics Review, July 30, 2001





Opto-IRAM

Overview of Technical Approach

Central Concepts:

- Reduced memory latency in multiprocessors with on-chip RAM and wide, parallel, optical inter-chip network
- Synergy of Computational-RAM concept with global, optical inter-PE netw'k

Computing and network architecture:

- Fast scalar processor; on-chip SDRAM; array of bit-serial PEs at sense amps
- Cache-coherent, parallel variant of *S-Connect* for low-latency shared memory

Advanced OE module concept:

- Free-space optics on-module; compliant guided-wave optics between modules
- Simple physical design with common beam specification at module periphery

Design automation:

- Conjoint optimization of bimodal free-space-optical/wire global interconnects

Modeling:

- Holistic model linking system performance to OE component characteristics

Optoelectronic Intelligent RAM


DARPA VLSI Photonics Review, July 30, 2001





Opto-IRAM Rationale

- Processor-memory latency requires implementing resource-intensive latency-hiding mechanisms.
- Bring DRAM onto the CPU: Intelligent RAM (IRAM) ...
- A single IRAM node does not have scalable memory.
- Network multiple chips together to expand memory ...
- Going off-chip for more memory forfeits the latency advantage of on-chip DRAM.
- Use the physical bandwidth and parallelism of VLSI Photonics program technologies, together with simple distributed memory protocols, to implement off-chip memory access with low latency ...



 Opto-IRAM Multiprocessors!

Optoelectronic Intelligent RAM

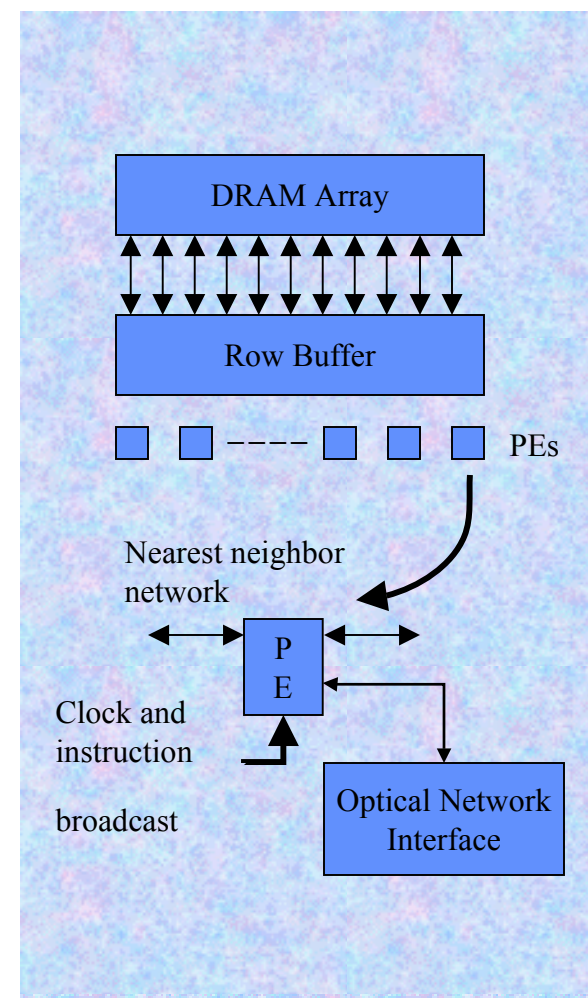
DARPA VLSI Photonics Review, July 30, 2001





Optical Computational RAM (*OCRAM*) Rationale

- Conventional memory system architecture masks the enormous memory bandwidth developed at the sense amplifiers.
- In addition to the scalar CPU, place simple (e.g., bit-serial) PEs in the DRAM banks . . .
- PEs need interconnection for data-parallel applications.
- Interconnect PEs into a mesh-like network . . .
- Inter-PE latency across mesh is large.
- Implement small-diameter inter-PE network using free-space optics . . .



Opto-CRAM concept!

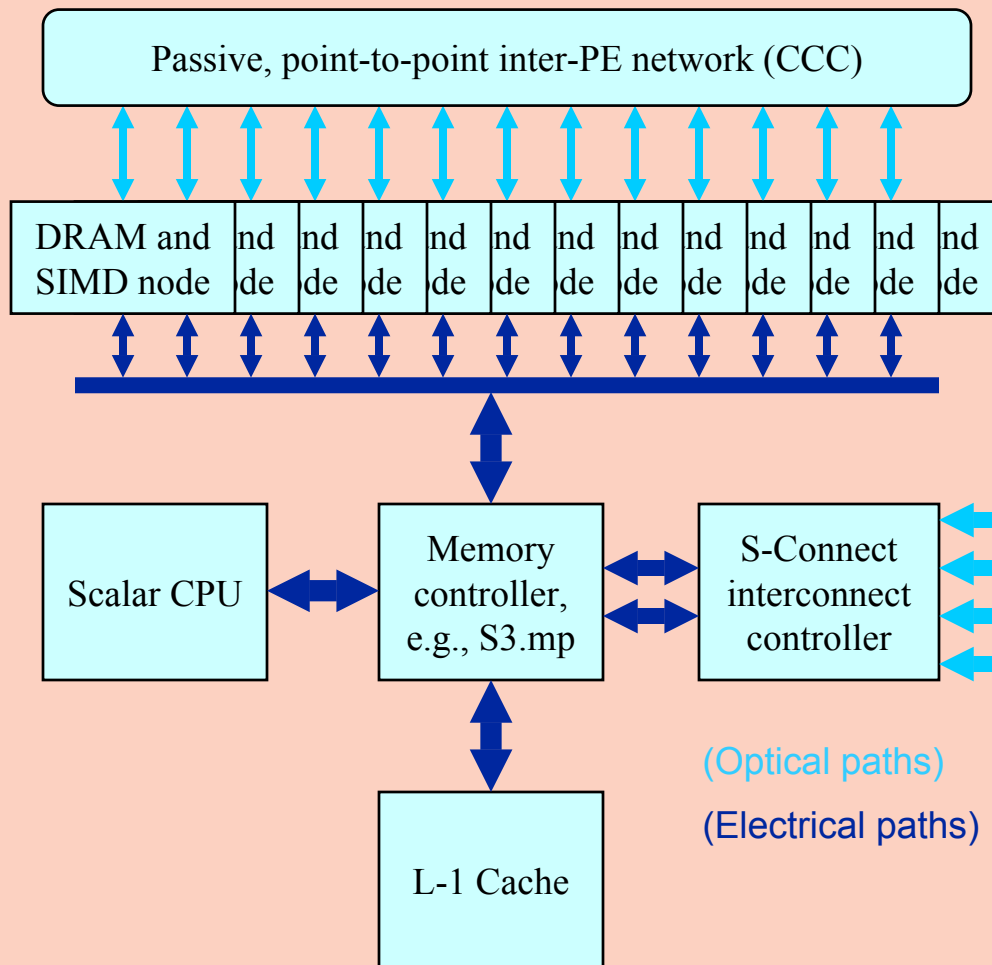
Optoelectronic Intelligent RAM

DARPA VLSI Photonics Review, July 30, 2001





Opto-IRAM/CRAM Node Overall Organization



- : CPU sees cache-coherent bus.
- : Computational RAM (CRAM) PEs exploit bandwidth at sense amplifiers.
- : Separate optical network facilitates global connectivity among PEs.
- : Low inter-chip cache line latency from parallel optics.

S-Connect:

- : Supports cache-coherent global address space
- : Scalable, point-to-point network
- : See <http://playground.sun.com/pub/S3.mp/s3mp.html>

Optoelectronic Intelligent RAM

DARPA VLSI Photonics Review, July 30, 2001





Unique Features of OCRAM

- **Decouple PE cycle from memory cycle**
 - Use double-buffering to overlap PE operation and memory access
=> multiple instruction streams, similar to pipelining
 - PE is simple and fast - can be clocked much faster than DRAM
 - Use destructive DRAM reads when possible - fast and low power
- **Balance compute bandwidth and memory bandwidth**
 - One PE for multiple DRAM columns
 - Many PE operations overlap each DRAM operation
 - Width constraints on PE are loosened - more efficient layout
- **Inter-PE communication bandwidth**
 - Solution: Use optical interconnect



Optoelectronic Intelligent RAM

DARPA VLSI Photonics Review, July 30, 2001



Sample OGRAM Microprogram (Add32) - 195 clock cycles

```
c ADD R2 0 31      ; init. Loop ctr.
p SIMD Y<-0        ; init. Carry in
c LOOP0 end R2 R3  ; R2 counts 31..0, R3 counts 0..31
c ADD R4 32 R3     ; calculate bit# of hiword[R3]
a BITSEL R4        ; M is now alias for hiword[R3]
p SIMD X<-M        ; save hiword[R3] in X
a BITSEL R3        ; M is now alias for loword[R3]
p SIMD M<=X^Y^M    ; calc. sum & store in loword[R3]
end:
p SIMD Y<=/M*(X+Y)+XY ; calc. Carry out & store in Y
```




Optoelectronic Intelligent RAM

DARPA VLSI Photonics Review, July 30, 2001



OCRAM System Simulator



Optoelectronic Intelligent RAM

DARPA VLSI Photonics Review, July 30, 2001





OCRAM Status and Goals

Status:

- OCRAM Instruction Set Definition Complete
- PE Preliminary Layout Complete
- Preliminary SPICE Simulations Performed (not extracted from layout)
- Microcode for several basic operations completed



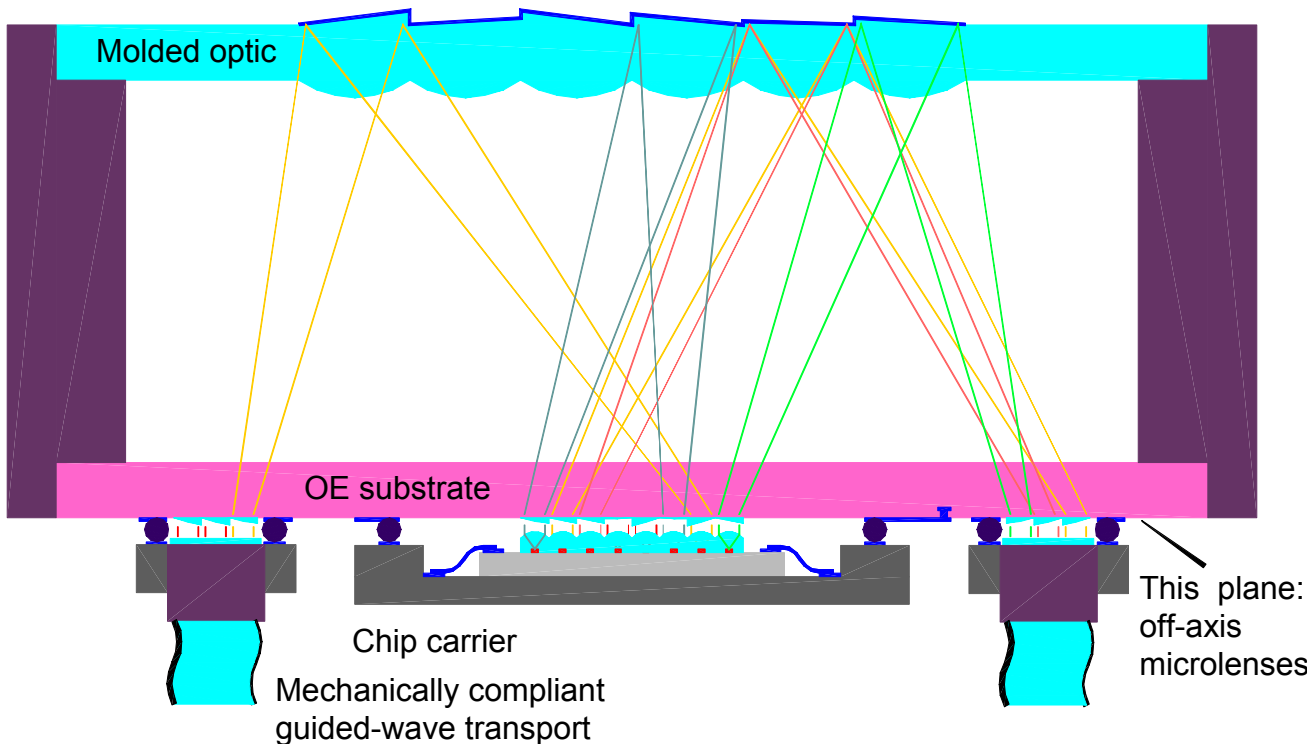
Optoelectronic Intelligent RAM

DARPA VLSI Photonics Review, July 30, 2001



Opto-IRAM Single-Chip Module

Scalable optical interconnect hierarchy



- : Free-space **intra**-chip, Guided-wave (compliant) **inter**-chip connections.
- : Single optical interface specification to OE module substrate.
- : Most critical alignment (microlenses to sources/detectors) is drawn within a highly controlled environment (VLSI/OE aggregate).
- : Free-space path between imaging bundle and chip provides routing and access to any chip ports.

Optoelectronic Intelligent RAM


DARPA VLSI Photonics Review, July 30, 2001






Regularity and Free-Space Optical Interconnection: CAD Opportunity

Fourier-plane system volume is sensitive to space-variance:





Some patterns with identical space-variance:



: Regularity matters, not path length.

Optimal placement and net assignment:



- : Shorter average wire length
- : Regular optical pattern
- : Physical compactness

Optoelectronic Intelligent RAM

DARPA VLSI Photonics Review, July 30, 2001





CAD for Partitioning/Placement

Problem statement:

- **Given:** Functional cells, netlist, and lattice of optical input and output ports.
- Place cells and assign nets to optical ports to minimize the cost function.

Cost function:

- Terms for : wire resources cost, : optical resources cost
- Wire resources cost term $\int \sigma + 2\sigma^2$ of wire length distribution
- Optical resources cost term \int space-variance of optical paths

Algorithm:

- Evaluate many standard-cell placements.
For each placement, evaluate best assignment of nets to optics;
For each assignment, evaluate the best routing of electrical nets:
 - \int To other nets,
 - \int To optical ports.
- Simulated annealing, genetic algorithms for stochastic optimization




Optoelectronic Intelligent RAM

DARPA VLSI Photonics Review, July 30, 2001



Partitioning/Placement Approach



Optoelectronic Intelligent RAM

DARPA VLSI Photonics Review, July 30, 2001





Optimization

- Simulated Annealing
- Genetic Optimization
- Cost Function tries to maximize utilization of optical bandwidth
- Layout compaction done after each “move”..

$$\text{Cost function} = \sum (\text{electrical wire length})^n + \left(\frac{K}{\text{utilization of optical bandwidth}} \right)^m$$

$$\text{Utilization} = (\text{Bandwidth utilized} / \text{Available bandwidth})$$








Results: Branch and Bound

No. of gates	% of wires converted to optical links	% reduction in longest wire length	% reduction of modified mean	No. of optical directions
625	24	45	30	69
2500	23	40	29	120
5000	26	47	33	170
10000	21	37	25	280

(a)

(b)

(c)




Optoelectronic Intelligent RAM

DARPA VLSI Photonics Review, July 30, 2001





Cost Reduction vs. No of Optical Directions



The graph for cost vs. optical directions with various number of modules


Optoelectronic Intelligent RAM

DARPA VLSI Photonics Review, July 30, 2001





Routing Without Optimization



Routing without optimization


Optoelectronic Intelligent RAM

DARPA VLSI Photonics Review, July 30, 2001





Optimized Optical Nets



Optical routing with optimization


Optoelectronic Intelligent RAM

DARPA VLSI Photonics Review, July 30, 2001





Optimized Electrical Nets



Electrical routing with optimization

Optoelectronic Intelligent RAM

DARPA VLSI Photonics Review, July 30, 2001





Optimization Results

No. of modules	% of wires converted to optical links	% reduction in longest wire length	% reduction of mean cost	No. of optical directions
9	25	39	22	246
16	44	44	28	437
25	57	52	40	567
36	60	69	45	598

Comparison with different number of modules

Optoelectronic Intelligent RAM

DARPA VLSI Photonics Review, July 30, 2001





Optimization With Diff Sensor Layouts

Arrange- ment	% of wires converted to optical links	% reduction in longest wire length	% reduction of mean cost	No. of optical directions
(a)	44	44	28	437
(b)	42	37	24	399
(c)	40	32	23	373

Comparison with different sensor distribution


Optoelectronic Intelligent RAM

DARPA VLSI Photonics Review, July 30, 2001





Optimization: Global vs. Local Interconnect



The graph for the reduction of longest wire vs. distance with 25 modules and 100 optical direction

Optoelectronic Intelligent RAM

DARPA VLSI Photonics Review, July 30, 2001





Optimization Results

		10	50	100	200	300	400	500	600
9	OD	9	49	96	175	208	249	246	246
	Total	2 %	10 %	18 %	30 %	34 %	38 %	37 %	38 %
	Longest	11 %	32 %	26 %	35 %	36 %	33 %	42 %	39 %
16	OD	9	49	99	195	285	345	389	437
	Total	2 %	9 %	19 %	33 %	44 %	51 %	55 %	58 %
	Longest	20 %	26 %	41 %	41 %	45 %	36 %	45 %	44 %
25	OD	9	49	99	199	299	395	486	567
	Total	2 %	10 %	19 %	34 %	48 %	58 %	67 %	73 %
	Longest	21 %	31 %	33 %	47 %	54 %	53 %	58 %	52 %
36	OD	9	49	99	199	299	399	498	598
	Total	2 %	10 %	19 %	34 %	49 %	60 %	69 %	78 %
	Longest	19 %	31 %	46 %	47 %	56 %	62 %	66 %	69 %
49	OD	9	49	99	199	299	399	499	596
	Total	2 %	10 %	19 %	35 %	61 %	61 %	70 %	78 %
	Longest	17 %	31 %	48 %	58 %	61 %	61 %	68 %	64 %

Optoelectronic Intelligent RAM

DARPA VLSI Photonics Review, July 30, 2001





Opto-IRAM

Summary of Accomplishments

Computing and network architecture

- OGRAM node architecture, programming model
- Cache-coherent, shared-memory multiprocessor architecture based on S-Connect

Holistic modeling and virtual prototyping

- Behavioral link modeling with Spice, Chatoyant

Optical interconnect architecture

- Advanced OE module physical design

Design automation

- New partition/placement tool for mixed-mode interconnect optimization



Optoelectronic Intelligent RAM

DARPA VLSI Photonics Review, July 30, 2001